

AD-A053 243

MASSACHUSETTS INST OF TECH CAMBRIDGE OPERATIONS RESE--ETC F/G 15/5
THE MULTI-PRODUCT PRODUCTION CYCLING PROBLEM: DEVELOPMENT OF HE--ETC(U)
MAR 78 S C GRAVES

UNCLASSIFIED

TR-146

N00014-75-C-0556

NL

1 OF 1
AD
A053243

TEC
146



END
DATE
FILED
5 - 78
DDC



AD No. —
DDC FILE COPY

AD A 053243

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 147 TR-1 Technical Report No. 146	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) ⑥ THE MULTI-PRODUCT PRODUCTION CYCLING PROBLEM: DEVELOPMENT OF HEURISTICS I	⑨	5. TYPE OF REPORT & PERIOD COVERED Technical Report March 1978
6. AUTHOR(S) ⑩ Stephen C. Graves	⑮	6. PERFORMING ORG. REPORT NUMBER
7. PERFORMING ORGANIZATION NAME AND ADDRESS M.I.T. Operations Research Center 77 Massachusetts Avenue Cambridge, MA 02139	8. CONTRACT OR GRANT NUMBER(S) N00014-75-C-0556	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NR 347-027
11. CONTROLLING OFFICE NAME AND ADDRESS O.R. Branch, ONR Navy Dept. 800 North Quincy Street Arlington, VA 22217	12. REPORT DATE Mar 1978	13. NUMBER OF PAGES 22 pages
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) ⑫ 269.	15. SECURITY CLASS. (of this report)	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Releasable without limitation on dissemination.	DISTRIBUTION STATEMENT A Approved for public release Distribution Unlimited	
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES	D D C RECEIVED APR 27 1978 A RECORDED	
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Multi-Product Production Scheduling Problem Markov Decision Problems Heuristic Decision Policies		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) See page ii.		

270 720

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

THE MULTI-PRODUCT PRODUCTION CYCLING PROBLEM:
DEVELOPMENT OF HEURISTICS

by

STEPHEN C. GRAVES

Technical Report No. 146

Work Performed Under
Contract N00014-75-C-0556, Office of Naval Research
Multilevel Logistics Organization Models
Project No. NR 347-027

Operations Research Center
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

March 1978

DISTRIBUTION STATEMENT A	
Approved for public release Distribution Unlimited	

ACCESSIONED FOR	
REF#	White Sackler
DDC	Off Sackler
UNANNOUNCED	
AUTHENTICATION	
BY	
DISTRIBUTION/AVAILABILITY CODED	
DISL	AVAIL AND/OR SPECIAL
A	

Reproduction in whole or in part is permitted for any purpose of the
United States Government.

FOREWORD

The Operations Research Center at the Massachusetts Institute of Technology is an interdepartmental activity devoted to graduate education and research in the field of operations research. The work of the Center is supported, in part, by government contracts and grants. The work reported herein was supported by the Office of Naval Research under Contract N00014-75-C-0556.

Richard C. Larson
Jeremy F. Shapiro

Co-Directors

Abstract

The multi-product production cycling problem is concerned with the determination of a production/inventory policy for a single capacitated production facility which is dedicated to producing a family of products. This paper studies this problem assuming stochastic demand. The one-product problem is formulated as a Markov decision problem which may be reasonably solved. For the multi-product problem, heuristic decision rules are proposed. In the context of an identical-product problem, we develop a heuristic decision policy which is based on the analysis of the one-product problem, and on two new notions: the composite product and the lead-time adjustment. This heuristic is then extended to the identical-cost problem, and the correlated demand problem. Arguments are presented for the generality of the identical-cost problem, and hence the generality of the proposed heuristic policy.

1. Introduction

The multi-product production cycling problem (MPCP) is concerned with determining production and inventory control policies for a group of products. Each of the products is produced or processed on a single facility. This production facility or machine has finite capacity and processes only one product at a time. The production for each of the products goes into inventory with demand for the product being serviced from this inventory. The MPCP is to devise a scheduling policy for the single production facility. The criteria for the scheduling policy is minimum cost per unit time, where cost consists of setup costs, inventory holding costs, and backorder costs.

Most of the literature on the MPCP has been confined to models which assume deterministic demand. Delporte and Thomas [1] and Silver [11] give comprehensive bibliographies for the MPCP with deterministic and constant demand rates. The work on stochastic demand problems has been either very preliminary in nature [3] or hampered due to extremely restrictive assumptions [5]. The deterministic results do not seem to be particularly helpful in solving the stochastic problem; unlike the classic deterministic inventory theory which provides the first step in the development of stochastic inventory theory, there has yet to be demonstrated a strong link between the deterministic and stochastic MPCP. The reason for this lack of congruence seems to lie in the complexity inherent in the problem: Any solution procedure for the MPCP must tradeoff lot-sizing decisions against capacity allocation decisions. For deterministic demand, the existing procedures focus primarily on lot-sizing; the capacity allocation decisions come into play only when capacity constraints are violated which force an adjustment in the lot-sizing decisions (e.g. see [2]). For stochastic demand, it is not possible to consider lot-sizing and capacity sequentially: A solution procedure which focuses on lot-sizing loses the flexibility needed to cope with the capacity

allocation decisions under unknown demand. Similarly, considering capacity allocation at the expense of lot-sizing will be too rigid to handle fluctuating demand and inventory levels. An effective solution procedure must be responsive to the limitations on scheduling due to finite capacity, and still be able to reflect the lot-sizing tradeoffs among inventory holding costs, backorder costs and setup costs. A solution procedure must simultaneously account for lot-sizing and capacity allocation.

The intent of this paper is to develop heuristic control policies for the stochastic demand MPCP. A subsequent paper [7] demonstrates the effectiveness of these heuristic policies. The specific problem to be considered is defined by the following assumptions:

1. Demand for each product is stochastic and is specified by a known stationary distribution function.
2. The form of the decision policy is periodic review.
3. At most one product may be produced in a period. Furthermore, the product must be produced for the entire duration of the period at the product's finite production rate.
4. Production in a period is available for filling demand orders at the end of the period; there is no production lead time.
5. A setup cost is incurred whenever a product is produced unless that product was produced (setup) in the immediately preceding period. That is, setups are carried over from prior periods. Whenever the facility is shutdown, any prior setup is lost and is not carried over.
6. Setup times are assumed to be zero.
7. Demand exceeding on-hand inventory is backordered.
8. The objective is minimization of the expected cost per unit time over the infinite horizon over all products. The components of the total cost are setup costs, inventory holding costs, and backorder costs.

These assumptions may seem overly restrictive, especially assumption 4 (zero lead time) and 6 (zero setup time). However, these two assumptions are made for ease of exposition and may be relaxed, as will be seen.

The remainder of the paper is divided into 4 sections. In the next section the one-product production cycling problem is examined. Based on the analyses of the one-product problem, Section 3 develops heuristics for a special form of the MPCP: the identical-product problem. In Section 4 these heuristics are extended to a more general class of MPCP. Section 5 gives a brief summary and discussion of the results of the paper.

2. The One-Product Problem

This section considers a special case of the MPCP: the one-product problem. This problem is not a very realistic problem. Our interest in this problem, however, stems from the fact that some theoretical analysis of the problem may be performed for this version of the MPCP. Insight gained from this analysis will form the basis for the study of the more complex and more realistic multi-product problem.

The one-product problem is first formulated as a Markov decision problem. A discussion is given for the solution of this problem, and for the form of the optimal policy. It is conjectured that the optimal policy is a two-critical-number policy. In addition, we define for the optimal policy a function which may be interpreted as a shadow price for the machine's capacity.

Formulation of One-Product Problem

Consider the one-product problem. The cost expression for this problem for a finite time horizon can be defined by backward recursion as follows:

$$(1) \quad C(I, k, t) = \min \{ G(I) + E_{\lambda} [C(I-\lambda, 0, t-1)]; \\ G(I+P) + [w(k)]K + E_{\lambda} [C(I+P-\lambda, 1, t-1)] \}$$

where

$C(I, k, t)$ = expected cost for remaining t periods assuming an optimal decision policy, given inventory level I and machine status k .

$C(I, k, t=0) = 0$ for all I, k

k = status of facility

$$k = \begin{cases} 0 & \text{if facility is not setup for production} \\ 1 & \text{if facility is setup for production} \end{cases}$$

$G(I)$ = expected backorder and inventory holding cost for current period, given net inventory level I available for servicing demand.

$E_{\lambda}[C(I-\lambda, k, t)]$ = expectation of the cost function taken over λ ,
the demand random variable.

= $\sum_{i=0}^{\infty} p(i) C(I-i, k, t)$ for $p(i)$ = probability that
demand (λ) equals i .

P = production rate measured in units per period.

K = setup cost

$w(k)$ = indicator function

= $1-k$

Thus, the expected cost from following the optimal decision policy for the remaining t periods is the minimum of: (1) the expected cost if we do not produce in the current period; (2) the expected cost if we do produce in the current period. In each case, the expected cost is the expected back-order and inventory holding cost for the current period, plus a setup charge, if applicable, plus the expected cost for the remaining $t-1$ periods. Note that the machine status (k) in period $t-1$ depends only on the current decision. Also, if the current decision is to produce, inventory is increased by P , and a setup charge is incurred only if the machine was previously not setup.

This formulation may be modified to model other versions of the one-product problem. It is trivial to extend the model if setups cannot be carried over or if the machine is always setup. The assumption of zero lead time may be relaxed to allow for a lead time equal to an integral number of periods; this extension is analogous to that given in [9] for a pure inventory model. The assumption of zero setup time can also be relaxed for a special case where not only is there a setup time but also a shutdown time. The length of the shutdown time must be the same as that for the setup time. The details of this modification are given in [6].

To determine optimal policies for the one-product problem, one can use
(1) to recursively compute the optimal production decisions for all values

of inventory I and remaining time t . Alternatively, if the time horizon for the production process is long, one is interested in finding the stationary steady-state optimal policy. To do this, the policy iteration method for sequential decision problems as developed by Howard [8] can be used. This method, however, requires a finite state space. The one-product problem has an infinite state space since the inventory level can range conceivably from $-\infty$ to $+\infty$. To apply this method, it is necessary to truncate this infinite state space to a finite state space by restricting the inventory level to range between some lower and upper inventory levels. By choosing these boundaries intelligently, the truncated problem will be an accurate representation of the original problem.

Form of Optimal Policy

It is conjectured that the form of the stationary steady-state optimal policy is a two-critical-number policy. That is, a policy may be characterized by (I^*, I^{**}) such that once inventory drops below I^* , the facility is turned on and is kept on until the inventory level reaches I^{**} ($I^{**} \geq I^*$) at which point the facility is shut down. This conjecture is very reasonable in the context of inventory theory; indeed it is difficult to fathom any other policy form being optimal. Nevertheless, this author has not been able to prove or disprove this conjecture.

It is difficult to pinpoint exactly why this conjecture has been so difficult to prove or disprove. The problem as formulated is more complex than many standard problems in the production/inventory literature in that it allows stochastic demand, non-zero setup cost, and, in particular, finite production capacity. The work of Scarf [10] in establishing the optimality of (S, s) policies for the pure inventory problem, and the work of Veinott [13] in showing the optimality of a single-critical-number policy

for a batch-ordering problem without fixed ordering costs, are directly related to the one-product problem, but have not been useful in proving the conjecture.

Despite the absence of a conclusive proof for the conjecture, there is strong supporting evidence for its validity. When the setup cost is zero ($K=0$) or when setups cannot be carried over, the optimal policy is characterized by a single critical number I^* . That is, if inventory is less than I^* , a decision to produce in the current period is made; otherwise there is no production. A proof is given in [6]. Another special case that can be analyzed is the one-product problem with positive setup cost but where the production quantity P is set to one unit and the product demand is always an integer number of units. It is proved in [6] that the optimal policy for this problem is a two-critical-number policy.

An additional indication of support for the conjecture is given by some results on the optimal control of a single-server queue. Sobel [12] considers the continuous control of a single queue for which there are start-up and shut-down costs for turning on and off the server, and there are general costs associated with the queue length. This problem is identical to a continuous review version of the one-product problem where the server corresponds to the production facility, and the queue length is the order backlog. Sobel shows that any reasonable policy will revert to a two-critical-number policy; hence, the optimal policy for the continuous review one-product problem is a two-critical-number policy. This is strong evidence for the conjecture's validity since continuous review policies are the limit of periodic review policies as the period lengths go to zero.

The Value Function

To gain insight into the behavior of the optimal policy for the one-product problem, it is useful to define the value function $V(I,t)$ as

$$(2) \quad V(I,t) = G(I) - G(I+P) + E_{\lambda}[C(I-\lambda, 0, t-1)] - E_{\lambda}[C(I+P-\lambda, 1, t-1)]$$

The optimal production policy is determined by this function. If the machine is setup ($k=1$), a decision to produce should be made if and only if $V(I,t)$ is positive; alternatively, if the machine is not setup ($k=0$), a decision to produce should be made if and only if $V(I,t)$ is greater than K , the setup cost. Note that if $V(I,t)$ is non-increasing in I , the optimal policy is a two-critical-number policy where $[I^*(t), I^{**}(t)]$ are defined by

$$(3) \quad V[I^*(t)+1, t] < K \leq V[I^*(t), t]$$

$$(4) \quad V[I^{**}(t)-1, t] > 0 \geq V[I^{**}(t), t]$$

$V(I,t)$ may be interpreted as the value of a producing versus non-producing production facility given an inventory level of I and t periods remaining; that is, a rational product manager would be willing to pay as much as $V(I,t)$ to turn on the machine in the current period, given that in subsequent periods it will cost K to setup the machine. Hence, if this value is greater than the actual cost of turning on the machine in the current period (K or 0, depending on the machine's status), then a decision to produce is made.

For the steady-state problem we define $V(I)$ to be the limit of $V(I,t)$ as t goes to infinity. $V(I)$ may be readily computed from the results of the policy iteration method. If $V(I)$ is non-increasing, the optimal steady-state policy is given by (I^*, I^{**}) defined similar to (3), (4).

3. The Identical-Product Problem

The multi-product problem can be formulated similar to (1) as a set of backward recursion equations. In theory, the multi-product problem can also be solved by means of the policy iteration method. However, the computation required by this method increases exponentially with the number of products, and becomes computationally infeasible for reasonably-sized problems. Consequently, we have concentrated on finding good heuristic policies which are easy to compute and implement. We use the identical-product problem to illustrate the development of these heuristics.

The identical-product problem is a special case of the MPCP in which all of the products have identical demand distributions, identical production rates, and identical cost structures, but with their demand realizations being independent across products.

A heuristic solution for this problem can be developed on the basis of the one-product analysis: Assume that each product is the only product produced by the machine; then the one-product analysis may be done to find a decision policy for each product. If we assume the conjecture on the form of the optimal policy is true, then each of the identical products is governed by a two-critical-number policy (I^*, I^{**}). Each product's policy may be executed as long as this does not conflict with the execution of another product's policy. A conflict will occur whenever more than one product wants to be produced in a given period. For example, consider the case where the machine is setup for product 1, product 1 has inventory (I_1) less than I^{**} , and product 2 has inventory (I_2) less than I^* ; here both products want the machine's capacity in the immediate period. The heuristic must resolve the conflict in a reasonable manner. One procedure might be to compare inventory levels relative to the respective critical

numbers (e.g. $I_1 - I^{**}$ or $I_2 - I^*$), and produce the product with the lower relative inventory. Another procedure would be to use the value function, and produce that product which values the machine the highest. Here $V(I_1)$ is compared against $V(I_2) - K$; note that if the product is not setup, the setup charge must be subtracted from the value function to give the effective value of the machine.

The Composite Product

The procedure based on the one-product analysis is reasonable; clearly if the one-product analysis indicates production for a product, then in the multi-product problem with competition for the available capacity, production is still desirable. However, a weakness in this heuristic is that it does not try to anticipate or plan for potential production conflicts. As an example, consider a two-product problem in which the machine is shutdown, and the inventory levels for both products are slightly above I^* . One-product analysis would indicate no production in the current period; however, it is likely in the subsequent period that both inventory levels will fall below I^* , and both products will want to use the production facility. Since only one product can be produced per period, the inventory for one of the products will continue to fall, incurring increasing costs, until it is able to wrest the production facility away from the other product.

This problem could be avoided with some foresight in the production scheduling procedure. By anticipating the potential competition for the production resource before it occurs, a decision to produce one of the products in the current period could be made.

To help anticipate possible conflict situations, the notion of a composite product is introduced. In the context of the two-product example, define the composite product to have the same cost parameters and the same

production rate as either of the products. Let the demand for the composite product be the total demand for the two separate products; that is, the demand distribution is the convolution of the demand distributions for the individual products. This composite product can now be used to recognize potential production conflicts. The composite product can be analyzed as if it were one product to find a production policy (I_c^*, I_c^{**}) . Now, with the composite inventory, $I_c = I_1 + I_2$, a decision whether or not to "produce" the composite product can be made by comparing the composite inventory with either I_c^* or I_c^{**} . A decision to "produce" the composite product would indicate that the composite inventory was not adequate for the composite demand, realizing that the production rate for the composite product is the same as that for either of the individual products. The composite product should be helpful in anticipating potential conflicts; when both individual product inventories are just adequate when considered separately, in composite the composite inventory should indicate the need for current production.

The notion of the composite product is extendable to the n identical-product problem. Here, $n-1$ composite products are defined; composite product j for $j=2,3,\dots,n$ is the product formed from the composition of j identical products. The definition of each composite product is completely analogous to that given in the two-product example; the cost structure and production rate for the j^{th} composite product are the same as the individual products, while its demand is the convolution of the demands for j products. For each composite product, the one-product analysis can be performed to find the production policy (I_j^*, I_j^{**}) , where the subscript j refers to the j^{th} composite product $j=2,3,\dots,n$. The composite products can now be incorporated into the heuristic procedure to give better detection and reaction to potential conflicts over the usage of the production facility.

The Lead-Time Adjustment

Another modification that can be made is a lead-time adjustment. The one-product analysis for the individual products assumes, by definition, that the machine is dedicated to the one product. Hence, implicit in the analysis is the assumption that whenever a production order triggers, the machine is idle and is able to start production immediately. For the multi-product problem, this is clearly not the case. When a product triggers, there is no guarantee that the machine is not committed to another product. Hence, due to the interaction of the other products, there is usually a delay between the triggering and the servicing of a production request. One approach to account for this delay is to estimate the expected size of the delay and to treat it as a deterministic lead time. Then the individual products may be modeled as one-product problems but with a lead time equal to the actual lead time plus the expected delay.

To determine this lead-time adjustment, consider an n identical-product system for which the expected machine utilization is U , $0 < U < 1$. This system is modeled as a birth-death process on the state space $\{i: i=0, 1, \dots, n\}$, where state $i=0$ corresponds to the machine being idle while state i ($i>0$) is the machine busy with $i-1$ products having triggered and waiting for production time. It is assumed that products trigger randomly at an unknown rate α . Thus, the transition rate from state i to state $i+1$ is $\alpha_i = (n-i)\alpha$, since at state i there are $(n-i)$ products that have not triggered. The transition rate from state i to $i-1$ is $\mu_i = \mu$, except for $\mu_0 = 0$, where μ is the service rate for completing a production run for any product. Assuming that the expected length of a product's production run corresponds to production of an economic order quantity (EOQ), then

$$(5) \quad \frac{1}{\mu} = \frac{\text{EOQ}}{P} ,$$

where P is the production rate. Note that this birth-death process is an approximate model of the n -product system. For the birth-death process it is necessary to assume constant rates of transition; for the actual n -product system, the transitions are governed by rates that are not constant but depend on other system variables such as current service time and time since last production run. The intent of the birth-death model is to have a tractable model from which an approximate understanding of the actual system behavior may be obtained.

The standard analysis (see Feller [4]) of the equations of motion for the birth-death process yields the following equations for the steady-state probabilities $\{e_i\}$:

$$(6) \quad e_i = \frac{\alpha^i}{\mu^i} \frac{n!}{(n-i)!} e_0 \quad i=0,1,2,\dots,n$$

To find e_i , it is necessary to know α , e_0 , and μ . Since the machine utilization is U , the probability of the machine being idle must be $1-U$; hence, $e_0 = 1-U$. The trigger rate α is found numerically from the fact that the probabilities sum to one ($\sum_{i=0}^n e_i = 1$); μ is given in (5).

The expected delay can now be expressed as

$$(7) \quad L^* = \sum_{i=1}^n \left(\frac{e_i}{1-e_0} \right) \left(\frac{i-1}{\mu} \right)$$

where $\left(\frac{i-1}{\mu} \right)$ is the expected delay for a triggered product given that there are $i-1$ products ahead of it.

The lead-time adjustment L^* is appropriate for the one-product analysis of the individual products for the n identical-product problem. An analogous lead-time adjustment for the j -product composite product, $j=2,3,\dots,n$, would also be desirable. However, any realistic stochastic model to approximate this adjustment is much more complex than the birth-death model used for the individual products, and is not easily tractable.

Statement of the Heuristic

We can now state a heuristic for the identical-product problem, incorporating the one-product analysis, the use of composite products and the lead-time adjustment. For convenience in stating the heuristic, suppose the products are arranged so that $I_1 \leq I_2 \leq \dots \leq I_n$, where I_j is the inventory level of product j . Let k denote the status of the machine and assume that the optimal policy for the j -product composite product is a two-critical-number policy (I_j^*, I_j^{**}) , for $j=2,3,\dots,n$. Let (I_1^*, I_1^{**}) denote the optimal policy for the individual products where the lead time has been adjusted by the expected delay given in (7). The heuristic is as follows:

1. Let $j=1$; if $k=0$, go to step 6. If $k=1$, go to step 3. If $k>1$, go to step 2.
2. (case where $j < k$) Set $\hat{I}_1 = \sum_{i=1}^j I_i$, $\hat{I}_2 = I_k + \sum_{i=1}^{j-1} I_i$. If $\hat{I}_1 \leq I_j^*$ and $\hat{I}_2 < I_j^{**}$, go to step 5. If $\hat{I}_1 > I_j^*$ and $\hat{I}_2 < I_j^{**}$, produce product k . If $\hat{I}_1 \leq I_j^*$ and $\hat{I}_2 \geq I_j^{**}$, produce product 1. Otherwise go to step 4.
3. (case where $j \geq k$) Set $\hat{I}_1 = \sum_{i=1}^j I_i$. If $\hat{I}_1 < I_j^{**}$, produce product k ; otherwise go to step 4.
4. If $j=n$, go to step 7; otherwise, set $j:=j+1$. If $j < k$, go to step 2; if $j \geq k$, go to step 3.
5. Decide to produce either product 1 or product k .
6. Set $\hat{I}_1 = \sum_{i=1}^j I_i$. If $\hat{I}_1 \leq I_j^*$, produce product 1. If $j=n$ go to step 7. Otherwise set $j:=j+1$ and repeat this step.
7. Shutdown the facility.

The heuristic concentrates on possible production of at most two products: product 1, the product with least inventory, and product k , the product currently setup. For each value of j the procedure examines at most two j -product composite products: the composite product consisting of the j products with least inventories (composite product $[j]$), and the composite

product of the $j-1$ products with least inventories plus product k (composite product $[j-1;k]$). The heuristic uses the one-product analysis for the j -product composite product to determine whether composite product $[j]$ or $[j-1;k]$ should be produced. Note that if product k is not one of the j least inventory products or if $k=0$, then composite product $[j]$ is not setup; composite product $[j-1;k]$ is always setup. A decision to produce the composite product $[j]$ and incur a setup charge, implies production of product l ; a decision to produce the composite product $[j-1;k]$ implies the continuation of production of product k . When the facility is currently shutdown ($k=0$), only the composite product $[j]$ is considered. When product k ($k \neq 0$) is among the j lowest inventory products, the two composite products are the same. If for all j ($j=1,2,\dots,n$), no composite product is chosen to be produced, then the facility is shutdown. If for some j , the one-product analysis indicates that both $[j]$ and $[j-1;k]$ should be produced, then there is a conflict over whether product l or product k should be produced (step 5). Two suggested procedures for resolving the conflict are to compare relative inventories ($\hat{I}_1 - I_j^*$ vs. $\hat{I}_2 - I_j^{**}$) or to compare effective values from the j -product composite product value function $[V_j(\hat{I}_1) - k]$ vs. $V_j(\hat{I}_2)$.

4. Extension to Non-Identical Products

The previous section has considered the identical-product problem, for which a heuristic procedure was developed. The current section considers the general multi-product problem where the products are not identical. A discussion is given for the applicability of the heuristics proposed for the identical-product problem to the general MPCP. It is seen here that the heuristics are not easily extended to the non-identical-product problem. However, for a particular class of non-identical-product problems where the products have identical cost structures but varying demand rates, the composite product heuristic can be applied. Furthermore, it is argued that these identical-cost problems are a very realistic class of problems to consider.

The General MPCP

For the general MPCP, cost structures, production rates, and demand distributions may vary across products. One heuristic, similar to that proposed for the identical-product problem, would be to treat each product as if it had a dedicated machine. The one-product analysis could be done for the individual products, and the heuristic would produce a product that had triggered at the start of a period. If more than one product triggers, then the heuristic would have to decide which triggered product to produce, either by comparing relative inventory levels or by using the products' value functions. As was discussed for the identical-product problem, this heuristic ignores the scheduling dependency across products and, is unable to anticipate potential conflicts for the machine's capacity.

The composite product heuristic was developed for the identical-product problem to help anticipate production conflicts. When the products are not identical, the use of composite products is limited due to problems of

defining and comparing the composite products. First, it is not clear how to define a composite product; when the products have differing costs and production rates, it is difficult to determine the most appropriate parameters for a composite product. One possible approach is to use a weighted average of the parameters for the individual products, but again, it is not clear how to specify the weights for the individual products. Second, even with a means for defining the composite products, a heuristic would be difficult to implement due to the large number of possible composite products. For n products, the number of possible j -product composite products is $\binom{n}{j}$. Not only must the results from the one-product analysis be stored for each of these composite products, but in each decision period a method is needed for searching over and comparing the composite products.

When the products are identical, these problems are easier to handle since all j -product composite products are identical and hence, only $n-1$ composite products need to be defined (i.e., $j=2,3,\dots,n$). Furthermore, for the identical-product problem, the decision procedure needs to consider at most two j -product composite products: the composite product with least inventory and the composite that is currently setup with least inventory.

The composite product heuristic may be useful for the general MPCP by limiting the number of composite products that are considered. For instance, the decision procedure, in addition to the individual one-product analysis, would only consider a composite product consisting of all products or only composite products from the most popular products (e.g. "A" items in an "ABC" breakdown of inventory).

It would be desirable for the general MPCP to include a lead-time adjustment to the one-product analysis of the individual products. Unfortunately, though, to find this adjustment requires the analysis of a multi-dimensional process which is not tractable. Hence, any lead-time adjustment

would have to be made empirically or by rule-of-thumb.

The Identical-Cost Problem

A special case of the general MPCP for which the composite-product heuristic is directly applicable is what is called the identical-cost problem. Here, for all products, inventory holding and backorder costs are linear, and proportional to production rates; the setup costs for all products are the same. That is, for $i=1,2,\dots,n$

$$(8) \quad h_i = h/p_i$$

$$(9) \quad b_i = b/p_i$$

$$(10) \quad K_i = K$$

where

h_i = inventory holding cost per unit per period for product i ,

b_i = backorder cost per unit per period for product i ,

p_i = production rate in units per period for product i ,

K_i = setup cost for product i ,

h, b, K = positive constants.

The demand distributions may vary across products. Now, if a unit of each product is redefined so that all products have the same production rates (e.g., rescale $p_i = 1$ for all products), then the products will have identical cost parameters.

This problem and its assumptions are very realistic. Consider the inventory holding and backorder costs; these costs, in general, are assumed to be linear with respect to the cost of a product. In a single machine environment the cost of a product consists of the cost of the raw material as input to the processing facility plus the value added to this input during the processing function. Assuming that all products require the same or

similar raw materials, then it is reasonable to suppose that the machine processes the raw materials at a relatively steady dollar rate, independent of product. Furthermore, the rate of value added by the machine should be approximately constant across products. If this were not true, then it would be argued that the machine capacity is not being used efficiently. Thus, the value of product processed by the machine is nearly constant over time regardless of the particular product. Hence, after rescaling the products to a common production rate, the products should have similar inventory holding and backorder costs. Setup costs generally reflect the labor cost for altering the setup of the machine. If all products require a similar amount of work for their setups, then it is reasonable to assume they have the same setup costs.

Given identical-cost products, composite products are easily defined; all composite products have inventory holding cost h , backorder cost b , setup cost K , and the common production rate, while the demand distribution for the composite product is just the convolution of the demand distribution for the individual products. However, since demand rates vary over the individual products, there are still $\binom{n}{j}$ j -product composite products to consider. When all products were identical, for each $j=1,2,\dots,n$ at most two j -product composite products were considered. Similarly, for identical-cost products, attention can be restricted to two j -product composite products, for $j=1,2,\dots,n$. For (I_1, I_2, \dots, I_n) the inventory status of the n products, suppose that the products are numbered so that $(I_i/d_i) \leq (I_{i+1}/d_{i+1})$ where d_i is the demand rate for product i . That is, the products are arranged in increasing order of inventory measured in periods of demand (or "days of supply"). Now, completely analogous to the identical-product analysis, the two j -product composite products for consideration are the composite product formed from products 1 to j , and

possibly, the composite product formed from products 1 to $j-1$ with product k where the machine is currently setup for product k . These two composite products are then compared by means of their one-product analysis.

Thus, the notion of composite products, which was developed for the identical-product problem, can be extended to the more realistic identical-cost problem. Again, though, there is no tractable model to quantify a lead-time adjustment. Nevertheless, in practice, it would be desirable to recognize the need for a lead-time adjustment to reflect the scheduling dependency across products. This might be done by arbitrary increases in the trigger levels for the products.

Correlated Demand

When the demand across products is not independent but is correlated, the proposed heuristics can still be applied. Consider the identical-product problem; for the identical-cost problem, the treatment is analogous to that given for the identical-product problem. When demand is correlated, the one-product analysis for the individual products is not altered; in this analysis, there is no way to reflect the demand dependency. However, the composite product heuristic is ideally suited to account for any demand dependency; the demand distribution for a composite product naturally includes any correlation between the individual products. For instance, if a two-product composite product consists of two identical products, each with normally-distributed demand with mean μ and variance σ^2 , and if the covariance of demand between the products is $\rho\sigma^2$, then the composite product has normally-distributed demand with mean $\mu_c = 2\mu$ and variance $\sigma_c^2 = 2\sigma^2(1+\rho)$.

5. Discussion and Summary

In this paper the MPCP has been defined and analyzed. First, the one-product problem was formulated as a Markov decision problem, and could be solved in this form. It was argued that the multi-product problem, in general, is too complex to solve optimally. Hence, heuristic decision procedures were proposed. In particular, heuristics were developed for the identical-product problem; these heuristics were then extended to the non-identical-product problem, in particular the identical-cost problem, and to the correlated-demand problem. In the next paper [7] these heuristics are tested over a range of sample problems by means of simulation in order to demonstrate their effectiveness.

References

1. Delporte, C. M., and L. J. Thomas, "Lot Sizing and Sequencing for N Products on One Facility", Management Science, Vol. 23, No. 10, June 1977, pp. 1070-1079.
2. Doll, C. L., and D. C. Whybark, "An Iterative Procedure for the Single-Machine Multi-Product Lot Scheduling Problem", Management Science, Vol. 20, No. 1, September 1973, pp. 50-55.
3. Evans, R. V., "Inventory Control of a Multiproduct System with a Limited Production Resource", Naval Research Logistics Quarterly, Vol. 14, No. 2, June 1967, pp. 173-184.
4. Feller, W., An Introduction to Probability Theory and Its Applications, Volume I, Third Edition, New York: John Wiley and Sons, Inc., 1968.
5. Goyal, S. K., "Lot Size Scheduling on a Single Machine for Stochastic Demand", Management Science, Vol. 19, No. 11, July 1973, pp. 1322-1325.
6. Graves, S. C., "The Multi-Product Production Cycling Problem", unpublished dissertation, University of Rochester, October 1977.
7. Graves, S. C., "The Multi-Product Production Cycling Problem: Testing of Heuristics", January 1978.
8. Howard, R. A., Dynamic Programming and Markov Processes, New York: John Wiley and Sons, Inc., 1960.
9. Karlin, S., and H. Scarf, "Inventory Models of the Arrow-Harris-Marschak Type with Time Lag", Chapter 10 in Arrow, K. J., S. Karlin, and H. Scarf (eds.), Studies in the Mathematical Theory of Inventory and Production, Stanford, CA: Stanford University Press, 1958.
10. Scarf, H., "The Optimality of (S,s) Policies in the Dynamic Inventory Problems", Math Methods in the Social Sciences, Arrow, Karlin, and Suppes (eds.), 1959, pp. 196-202.
11. Silver, E. A., "A Simple Method of Determining Order Quantities in Joint Replenishment under Deterministic Demand", Management Science, Vol. 22, No. 12, August 1976, pp. 1351-1361.
12. Sobel, M. J., "Optimal Average-Cost Policy for a Queue with Start-Up and Shut-Down Costs", Operations Research, Vol. 17, No. 1, January-February 1969, pp. 145-162.
13. Veinott, A. F., "The Optimal Inventory Policy for Batch Ordering", Operations Research, Vol. 13, No. 3, May-June 1965, pp. 424-432.